REMARKS

The Examiner is thanked for the performance of a thorough search.

SPECIFICATION

In the specification, paragraphs [0005], [0009], and [0029] have been amended to provide typographical corrections. Specifically, in paragraph [0005] "associated" is corrected to read "associates", in paragraph [0009] "identify" is corrected to read "identity", and in paragraph [0029] "may used" is corrected to read "may be used" and "identify" is corrected to read "identity".

The Applicant respectfully submits that the correction of these typographical errors is evident from the context of the paragraphs themselves. No new matter is added.

STATUS OF CLAIMS

Claims 1-84 have been amended.

No claims have been cancelled, added, or withdrawn.

Claims 1-84 are currently pending in the application.

SUMMARY OF THE REJECTIONS/OBJECTIONS

Claim 7 has been objected to because of an informality.

Claims 1, 22, 43, and 64 have been rejected under 35 U.S.C. § 102(b) as allegedly anticipated by U.S. Patent Number 6,763,384 B1 issued to Gupta et al. (" *Gupta* ").

Claims 1-4, 6, 9, 11-14, 18, 22-25, 27, 30, 32-35, 39, 43-46, 48, 51, 53-56, 60, 64-67, 69, 72, 74-77, and 81 have been rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over U.S. Patent Application Publication No. 2003/0105856 A1 of Tse et al. (" *Tse* ") in view of U.S. Patent Application Publication No. 2002/0156882 A1 of Natarajan et al. (" *Natarajan* ").

Claims 5, 26, 47, and 68 have been rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over *Tse* in view of *Natarajan* and in further view of U.S. Patent Number 6,944,659 B2 issued to Taggart et al. (" *Taggart* ").

Claims 7, 8, 15, 28, 29, 36, 45, 49, 50, 57, 70, 71, and 78 have been rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over *Tse* in view of *Natarajan* and in further view of U.S. Patent Application Publication No. 2003/0177183 A1 of Cabrera et al. (" *Cabrera* ").

Claims 10-12, 31-33, 52-54, and 73-75 have been rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over *Tse* in view of *Natarajan* and in further view of the non-patent reference titled "SNMP Alarms and MIB Module" of Perkins (" *Perkins* ").

Claims 16, 37, 58, and 79 have been rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over *Tse* in view of *Natarajan* and in further view of U.S. Patent Number 6,425,008 B1 issued to Lecheler et al. (" *Lecheler* ").

Claims 17, 38, 59, and 80 have been rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over *Tse* in view of *Natarajan* and in further view of U.S. Patent Application Publication No. 2004/0213224 A1 of Goudreau et al. (" *Goudreau* ").

Claims 19-21, 40-42, 61-63, and 82-84 have been rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over *Tse* in view of *Natarajan* and in further view of U.S. Patent Application Publication No. 2003/0110398 A1 of Dacier et al. (" *Dacier* ").

The rejections are respectfully traversed.

RESPONSE TO REJECTIONS NOT BASED ON THE PRIOR ART

Claim 7 has been objected to due to an informality. The Office Action correctly observes that "more edge routers" has been written as "moreedge routers" in Claim 7.

The Applicant has amended Claim 7 to remove "one or moreedge routers" and insert "an edge router". The Applicant respectfully submits that this amendment to Claim 7 traverse the objection to Claim 7.

RESPONSE TO REJECTIONS BASED ON THE PRIOR ART

A. CLAIM 1 AS REJECTED BASED ON *GUPTA*

(1) INTRODUCTION TO CLAIM 1

Claim 1 features:

"A method for communicating an alarm in a computer network, comprising:

a **network device** ***detecting*** an event ***within*** the **network device** on the computer network, wherein the **network device** is included in a **particular site** in a plurality of sites, and wherein the event results from a ***change in operation of*** the **network device**;

in response to detecting the event, the **network device** generating and propagating an alarm to an **alarm identification component** that is ***hosted within*** the **network device**;

the alarm identification component augmenting the alarm with identification information to create an augmented alarm, wherein the **identification information** ***uniquely identifies*** the **particular site** among the plurality of sites; and

transmitting the augmented alarm to a network operations center for the computer network, wherein the **network operations center** is ***external*** to the **particular site** and the network operations center processes alarms for each site in the plurality of sites." (Emphasis added.)

Thus, Claim 1 features a network device that hosts the alarm identification component (AIC). The network device detects an event within itself due to a change in operation of the network device, and the alarm for the event is propagated to the AIC that augments the alarm, and then the alarm is transmitted to the network operations center (NOC). The AIC augments the alarm with identification information that uniquely identifies the particular site from the rest of the sites in the plurality of sites. At the NOC, which is external to the particular site that includes the network device, alarms are processed from the plurality of sites.

As a result of the approach of Claim 1, one or more potential benefits may be realized. For example, the use of private IP addresses and network address translation, which can affect

the IP address associated with the alarm, does not change the identification information. Therefore, the NOC can determine from which site of the plurality of sites the augmented alarm is transmitted based on the identification information used to augment the alarm, thereby avoiding the situation of having duplicate IP addresses that can arise with private IP addressing or having alarms identified with global IP addresses that may change from connection to connection, such as when using dynamic network address translation. Also, through the use of the identification information that uniquely identifies the particular site, the NOC does not need to track addresses or identities of individual network devices within each of the sites in the plurality of sites and then use those individual addresses or identities that are included in individual alarms to determine which site an alarm for a given network device belongs. Note that these potential benefits of the approach of Claim 1 are representative examples only, and Claim 1 and the remaining claims are not limited to implementations that provide all or even any of these example potential benefits that are discussed herein.

The amendments to Claim 1 are fully supported by the application, and no new matter is added. For example, FIG. 1C illustrates an embodiment in which each of devices 154A, 154B, 154C, and 154D of site A 150 include AIC 140, and similarly each of devices 156A, 156B, 156C, and 156D of site B 160 include AIC 140. This is in contrast to other embodiments, such as in FIG. 1B, in which AIC 140 is not hosted in devices 154A, 154B, 154C, and 154D of site A 150 and devices 156A, 156B, 156C, and 156D of site B 160, but rather AIC 140 is hosted separate from the devices, such as in a dedicated device for supporting AIC 140 or as part of an existing device, such as a router.

As described in the Application, in the embodiment illustrated FIG. 1C, if the event is associated with device 154A, then device 154A detects the event. (Paragraph [0032].) For example, the event may be when the network bandwidth available to the network device falls below a specified level or if the network device experiences a condition, such as the utilization of a processor on the network device is over 90%, both of which are examples of changes in operation of the network device. (Paragraph [0028].) In FIG. 1C, AIC 140 is deployed in individual network devices within the two sites being monitored by the MSP110, and as a result, when an alarm is generated by one of the network devices of FIG. 1C, then the alarm is propagated to AIC 140 that is hosted by that network device. (Paragraph [0040].) As illustrated in the embodiment of FIG. 2, AIC 140 augments the alarm with identification

information to create the augmented alarm, and in one embodiment, the identification information uniquely identifies the particular site from among the other sites in the plurality of sites. (Paragraph [0042].) Finally, as illustrated in FIG. 1A, 1B, and 1C, NOC 120 is external to site A 150 and site B 160 that include the network devices.

As a result of augmenting the alarm with the identification information that uniquely identifies the particular site, one or more of the problems discussed in the Background section of the Application may be addressed in a particular implementation, although the claims are not limited to implementations that address these particular problems.

For example, with the approach of Claim 1, the MSP need not rely upon an IP address, which is included with the alarm as a means to identify the source of the alarm, in order to determine from which site the alarm originated. As a result, other techniques that involve IP addresses, such as private IP addressing and network address translation, which can prevent the MSP from matching the IP address included with an alarm to a specific device, do not present a problem for the MSP. This is due to the MSP using the identification information to identify the site from which the alarm was transmitted, in contrast to using an alarm's source IP address that may be duplicated (e.g., for private IP addressing, two devices in different sites may have the same IP address) or an alarm's source IP address that may be masked or changed (e.g., network address translation that uses global IP addressing that shields the internal IP addresses to devices outside the site).

Also, as a result of the network device detecting the event itself and also hosting the alarm identification component that augments the alarm, the alarm with the identification information is transmitted to the network operations center without the need to have another device receive the alarm, augment the alarm, and then send the augmented alarm to the NOC.

(2) INTRODUCTORY DISCUSSION OF *GUPTA*

In contrast to the approach of Claim 1, *Gupta* discloses an approach to allow users to enroll to receive desired messages that includes the use of a receiving address identifier for the user. When a message monitor detects an event, the notification server determines the appropriate recipient and sends the notification to the user at the receiving address identifier. (Abstract).

For example, in Figure 3, *Gupta* illustrates a notification server 30 that communicates via the Internet 12 with both clients, such as "client 1" 114, "client 2" 115, "client 3" 116, and "client n" 118, and application servers, such as "application server 1" 20, "application server 2" 22, and "application server k" 24. Figures 4 and 5 show variations of Figure 3, such as in Figure 4 that includes "broker 1" 200 and "broker 2" 202 that act as intermediaries between notification server 30 and the clients, and such as in Figure 5 that illustrates clients behind a firewall and the use of socks server 220 to facilitate communication with the clients behind the firewall.

In particular, *Gupta* explains that a message monitor is running on each of application servers 20, 22, and 24, and that the message monitor "detects the occurrence of messages (i.e. 'events'), captures these messages and sends them to the notification server 30." (Col. 5, lines 31-34). When notification server 30 receives messages from the message monitor, "the notification server 30 determines the intended recipients of the messages using the databank of messages that the clients 110-118 wish to receive…Once the notification server 30 has received messages and identified the intended recipients, it will generate notification that are sent to the intended recipients that are currently on-line." (Col. 6, lines 11-24.) Thus, the notification server acts as a matching and forwarding server for messages identified by the message monitor by matching the identified messages from the message monitor to recipients that have signed up to receive those particular messages and are online.

Note that in the approach of *Gupta*, the "events" are merely messages generated by the application servers, that the message monitors capture those messages for forwarding to the notification server, and then that the notification server forwards the messages to clients that desire the messages based on a databank that maps clients to desired messages. Thus, a message ultimately sent to client has been handled by different servers: the application server that generated the message (e.g., the server where the "event" occurred) and upon which the message monitor is executing to capture the message (e.g., the message monitor detects the event), and the notification server that identifies the client as wanting that message based on the databank, and then forwards the message to the client.

(3) THE OFFICE ACTION'S REJECTION OF CLAIM 1 BASED ON *GUPTA*

The Office Action cites *Gupta* in rejecting Claim 1, referring to "the reasons given in the International Search Report" (ISR), which the Applicant understands to be a reference to the corresponding PCT application, PCT/US04/35867, that claims priority from the present Application and that was filed with the same claims as originally filed in the present Application. More specifically, the Applicant understands that the Office Action is referring to the "Written Opinion of the International Searching Authority" included in the International Preliminary Report on Patentability (IPRP) mailed on May 11, 2006 for PCT application PCT/US04/35867. And the Applicant interprets the Office Action's reference to "the reasons given in the International Search Report" (ISR) regarding *Gupta* as referring to Box. No. V.2. of the Written Opinion portion of the IRPR that refers to Claims 1, 12, 14, and 16, which the Applicant observes correspond to Claims 1, 22, 43, and 64 of the present Application. The Applicant notes that the Written Opinion alleges that Claims 1, 12, 14, and 16 of PCT application PCT/US04/35867 lack novelty under PCT Article 33(2) as being anticipated by *Gupta*, and the Written Opinion cites "column 2, line 64 – column 3, line 11" of *Gupta*.

Thus, the Applicant will address this citation of *Gupta* below as being the basis for the Office Action's rejections based on *Gupta* for Claims 1, 22, 43, and 64 of the present Application. However, if any of the above information regarding the Applicant's understanding of the basis for the Office Action's rejection of Claims 1, 22, 43, and 64 based on *Gupta* is incorrect, the Applicant respectfully requests that the next communication clarify any such misunderstandings by the Applicant.

The cited portion of *Gupta* is essentially a method claim in narrative form for notifying client processes of the occurrence of an event, presumably of the claims in *Gupta* as originally filed. Specifically, the method being described comprises the steps of: the client processes registering with a server both (a) events of interest and (b) the clients' respective address identifiers; the server detecting the occurrence of the event (presumably by being informed of the event by the message monitor on an application server that generates the message referred to by *Gupta* as an "event"); the server identifying which client processes are both (a) interested in notification of the event and (b) active; and the server causing a real-time connection to the interested and active client to transmit the notification to the interested and

active clients. (Col. 2, line 64 – Col. 3, line 11.) However, Claim 1, as amended above, differs from this description in Gupta in at least the following ways.

First, Claim 1 as amended above, features that "the event results from a change in operation of the network device." In contrast, *Gupta's* "events" are merely the occurrence of messages generated by the application servers (Col. 5, lines 31-34), and thus an "event" in *Gupta* does not result from a change in operation of the application server.

Second, Claim 1 features that the network device at which the event occurs is also the same network device that propagates the alarm to the alarm identification component that is also hosted by the same network device, and then the augmented alarm is transmitted to a network operations center. Thus, Claim 1 involves one device, the network device. In contrast, *Gupta* involves two devices that are separated by the Internet, one at which the event occurs and another at which the notification is generated to be sent to the client. Specifically, the messages in *Gupta* are generated by an application server, such as application servers 20-24, a message monitor on the application server identifies the messages and then transmits the messages to the notification server 30, which then performs the steps of the narrative method claim being cited in *Gupta*.

Thus, the "event" in *Gupta*, namely the message, occurs in an application server, which as illustrated by Figures 1, 3, 4, and 5, is separate from the notification server that identifies which active clients are enrolled to receive the particular message. In other words, the event occurs at one server (e.g., the application server) and then the forwarding of message as a notification to the client occurs at another server (e.g., the notification server) that is separate and different from the server at which the event occurred.

Also note that as illustrated in Figures 1, 3, 4, and 5, the Internet interposed between the application server and the notification server. Thus, the application server and the notification server are not only separate servers, but the application server and the notification server are not even part of the same site. This is fundamentally different than Claim 1 in which the network device is where the event occurs, the network device detects the event, and the network device generates and propagates the alarm to the alarm identification component that is hosted by the network device. Thus, in Claim 1, these steps occur not only within the particular site, but at the network device that is included in the particular site, whereas in *Gupta*, the event occurs and is detected at one server that is part of one site and then the

notification server that identifies clients to send the notification to is another server in another site.

Third, in Claim 1, the "identification information" that augments the alarm "uniquely identifies the particular site among the plurality of sites." In contrast, there is nothing in the cited portion of *Gupta* or any other portion of *Gupta* that the Applicant has reviewed that describes the notification server either modifying the messages or including in the notification anything that uniquely identifies a site that includes the application server from which the message originated (e.g., the "event"). Rather, the Applicant understands *Gupta* as disclosing that the notification server matches the messages from application servers to the active clients that are interested in the messages and then sends the messages to those active and interested clients. (Col. 6, lines 54-61.) Thus, the Applicant sees nothing in *Gupta* that would correspond to "identification that uniquely identifies the particular site among the plurality of sites," little less that such information is being used to augment an alarm.

Fourth, Claim 1 features "transmitting the augmented alarm to a network operations center for the computer network" and "the network operations center processes alarms for each site in the plurality of sites." The Applicant sees nothing in the cited portion of *Gupta* or any other portion of *Gupta* that discloses anything akin to such a network operations center as featured in Claim 1. As defined in the present Application, a network operations center (NOC) is a "software or hardware component that allows a MSP to receive and process alarms," (paragraph [0021]), and a "managed service provider (...MSP) is an entity, usually a business, which manages one or more computer networks that are each used by other entities (usually customers of the MSP)" (paragraph [0002]).

The Applicant respectfully submits that the clients in *Gupta* are not NOC's that are part of MSP's since the *Gupta's* client's do not manage other computer networks, but instead just receive notifications of messages from application servers that each client has registered an interest in receiving. Furthermore, the different clients in *Gupta* are understood to be part of separate networks themselves, and it is also understood that multiple clients would typically receive notification of a particular event from a particular server with *Gupta's* approach, while some clients would receive no notifications from some servers. Thus, there is nothing in *Gupta* to suggest that one particular client would receive notification of all events from all

application servers, little less that a client in *Gupta* "processes alarms for each site in the plurality of sites," as in Claim 1.

(4) CONCLUSION OF DISCUSSION OF CLAIM 1 AND *GUPTA*

Because *Gupta* fails to disclose, teach, suggest, or in any way render obvious "the event results from a change in operation of the network device," "a network device detecting an event within the network device…the network device generating and propagating an alarm to an alarm identification component that is hosted within the network device," "the alarm identification component augmenting the alarm with identification information to result in creating an augmented alarm," "the identification information uniquely identifies the particular site among the plurality of sites," "transmitting the augmented alarm to a network operations center for the computer network," and "the network operations center processes alarms for each site in the plurality of sites," the Applicant respectfully submits that, for at least the reasons stated above, Claim 1 is allowable over the art of record and is in condition for allowance.

B. CLAIM 1 AS REJECTED BASED ON *TSE* AND *NATARAJAN*

(1) INTRODUCTORY DISCUSSION OF *TSE*

In contrast to the approach of Claim 1, *Tse* discloses an approach for uniquely identifying an alarm notification that includes an alarm identifier field as a part of the notification and that the alarm identifier field has a "path" portion that identifies the series of nodes through which the alarm is relayed to the recipient. As a result, when acknowledging the alarm, the path portion of the alarm identifier field is used with the acknowledgement message to allow intermediate nodes to follow the path back to where the alarm notification originated. (Abstract.)

In particular, *Tse* illustrates in Figure 4A the structure of an alarm notification message 100 (paragraph [0037]) that includes an alarm identifier field 104 that in turn includes a path portion 108 (discussed further below) and an identifier portion 110 that comprise a string that identifies the alarm 100. (Paragraph [0042].) *Tse* explains that path portion 108 comprises a series of members, wherein each member is an identifier of an alarm supplier that handles the alarm 100. (Paragraph [0042].) *Tse* also explains that the alarm

identifier portion 110 that identifies alarm 100 is preferably unique for the alarm supplier that introduced the alarm to the network. (Paragraph [0042].) Alarm 100 also includes a system distinguished name field 102 that identifies the last node of the management system that carried the alarm notification, and as a result, if two nodes handle the alarm, system distinguished name field 102 identifies the second node to handle the alarm. (Paragraph [0042].) Finally, alarm 100 includes an alarm attribute field 106 that carries the alarm notification payload.

As a specific example, Figure 4b in Tse shows that alarm identifier field 104 includes path portion 108 that contains the values "0505.0010" for which "0505" identifies Member I and "0010" identifies Member II, thus identifying the two members through which alarm 100 of Figure 4B has passed. (Paragraph [0042].) Alarm identifier portion 110 of alarm identifier field 104 contains "9910," which is the alarm identifier assigned by Member II that has the identifier "0010." Because Member I was the last member to handle alarm 100, the system distinguished name field 102 contains "0505" that identifies Member I.

Thus, while *Tse* is altering the contents of alarm 100, the alterations are to include identifiers of the nodes through which the alarm passes, such that the path traversed by alarm 100 can be followed back step by step, such as when sending an acknowledgement of the alarm. However, such node identifiers only identify the individual nodes, and thus the node identifies do not identify the network or site in which each node resides. When the alarm in *Tse's* approach is generated by a node, the node only appends that node's identifier to the path portion 108 of the alarm identifier field 104, and the node's identifier only identifies the node, not the site that includes the node. The addition of other node identifiers to the path portion 108 merely identifies those particular nodes, not the sites that contain the nodes.

Because the alarm in *Tse's* approach passes through the monitored network 12 to the AMS 10 (see Figure 2), the path portion 108 would include node identifiers for nodes within both monitored network 12 and AMS 10. But neither the path portion 108 as a whole or the individual node identifiers for the nodes through which the alarm passes identify either the monitored network 12 or the AMS 10. And the system distinguished name field 102 merely reflects the node identifier of the last node to handle the alarm.

(2) THE OFFICE ACTION'S CITATIONS TO *TSE* IN REJECTING CLAIM 1

The Office Action states that *Tse* discloses "detecting an event associated with a device or any component thereof on the computer network ([0016, 0042]), in response to detecting the event, propagating an alarm to an alarm identification component ([0042]); at the alarm identification component, augmenting the alarm with identification information to result in creating an augmented alarm ([0042]); and transmitting the augmented alarm to a network operations center for the computer network ([0042-0043])."

As a preliminary administrative matter, the citations to *Tse* do not indicate what within those citations are being taken as corresponding to the particular features of Claim 1. For example, Claim 1 as amended above includes (1) a network device and (2) an alarm identification component hosted on the network device, yet the Office Action's citations to two detailed paragraphs in *Tse* do not allow the Applicant to determine what the Office Action is relying upon as disclosing either the network device or the alarm identification component that is hosted on the network device. In particular, paragraph [0042] merely discusses the structure of an alarm notification message, which can include identifiers for each alarm supplier that handles the alarm, but it is unclear what in paragraph [0042]. And in paragraph [0043], several different types of entities, such as alarm reporters, alarm collectors, and alarm consumers are referred to, yet it is unclear what, if any, of those entities the Office Action is relying upon as corresponding to the network device and the alarm identification component of Claim 1.

Therefore, the Applicant respectfully requests that any future Office Action explain what particular features of the cited reference is being relied upon as disclosing the features of Claim 1, and in particular, what is being relied upon as corresponding to the network device, the alarm identification component, and the identification information. Nevertheless, the Applicant has responded to the cited portions of *Tse* from the Office Action's rejections below and explained why the cited portions in *Tse* do not disclose many of the features of Claim 1 for which those portions of *Tse* are cited.

Claim 1 as amended above, features "the network device detecting an event within the network device," "the network device propagating an alarm to an alarm identification component that is hosted within the network device," and "the alarm identification component augmenting the alarm with identification information resulting in creating an augmented

alarm." Thus, the network device both (a) detects the event and (b) hosts the alarm identification component that augments the alarm with the identification information

Paragraph [0016] of *Tse*, which is cited by the Office Action in regards to the "detecting" step of Claim 1, explains that there is a new structure for uniquely identifying an alarm notification in which the node that creates the alarm and any nodes through which the alarm passes are identified through an alarm identifier field, thereby allowing an acknowledgement can be sent back along the same path that the alarm followed. However, this say nothing about a network device that both (a) detects an event and (b) hosts an alarm identification component. The passing of the alarm to other nodes involves separate devices, not the same device.

Paragraph [0042] describes the details of structure of an alarm notification message, which includes the following as illustrated in Figure 4A. However, a description of the details of the alarm notification message 100 says nothing about the detection of the event that resulted in the alarm, little less that such an event is detected by the device in which the event occurs or that the device propagates the alarm to "an alarm identification component that is hosted within the network device."

In addition, paragraph [0043] describes an example in which "an event, such as for example, a malfunction of a module of AR_1 202 occurs, action 2120, triggering a transmission of an alarm notification 212. According to the given implementation, the alarm notification 212 is transmitted to $AV_1$206 through AC_1 204." Thus, it appears to the Applicant that the Office Action is relying upon the alarm reporter 202 as disclosing the detection of an event.

However, in paragraph [0005], *Tse* explains with reference to Figure 1 that the "managed network 12 comprises one or more alarm reporters 14_i, such as for example the PC 14_1, the server 14_2, the door access device 14_3, etc. When abnormal conditions occur (such as for example a malfunction of the PC 14_1, or an open condition of device 14_3), alarm reporters 14_i issue alarm notifications 16, which are collected by the AMS 10 and relayed to one or more of the alarm viewers..." (Paragraph [0005].) As indicated in this description, the alarm reporters in *Tse's* approach are the devices themselves that generate the alarms, but there is nothing in this passage of *Tse* or any other portion of *Tse* about the devices propagating the alarm to something else within the device that augments the alarm with additional

information. The only thing like "augmentation" occurring within *Tse's* approach is the appending of node identifiers by each node that subsequently handles the alarm along the path to the alarm viewer, but those nodes are clearly separate from the alarm reporter that initially generated the alarm.

In addition, and assuming merely for the moment for argument's sake, that *Tse* does disclose the steps of "detecting," "propagating," and "augmenting" as featured in Claim 1, the Applicant fails to see anything within the cited portions of *Tse* or elsewhere within *Tse* that corresponds to "identification information uniquely identifies the particular site among the plurality of sites," as featured in Claim 1. The identifiers described in *Tse*, such as in paragraph [0042] and Figure 4A, are an alarm identifier portion 110 (of alarm identifier field 104) that identifies the alarm, a path portion (also of alarm identifier field 104) that comprises a series of identifiers of the members of an alarm supplier that handled the alarm, and a system distinguished name field 102 that identifies the last node of the management system that carried the alarm notification 100. All of these identifiers merely identify a node that handled the message, except for the alarm identifier portion 110 that identifies the alarm itself. But none of those identifiers "uniquely identifies the particular site among the plurality of sites," as in Claim 1.

Thus, the Applicant respectfully submits that contrary to the assertions of the Office Action, *Tse* fails to disclose, teach, suggest, or render obvious "the network device detecting an event within the network device," "the network device propagating an alarm to an alarm identification component that is hosted within the network device," "the alarm identification component augmenting the alarm with identification information resulting in creating an augmented alarm," and "identification information uniquely identifies the particular site among the plurality of sites," as featured in Claim 1.

(3) INTRODUCTORY DISCUSSION OF *NATARAJAN*

In contrast to the approach of Claim 1, *Natarajan* discloses an approach for identifying the source of an event in a computer network through the use of an "identifier tag" that "uniquely identifies at least one ***collection computer*** monitoring the event." A management computer receives information from the collection computer that includes the identifier tag,

and then the management computer ***derives*** an identification of each collection computer from the identifier tag. (Abstract; paragraph [0018]; emphasis added.)

In particular, *Natarajan* describes that the collection computers can be collection stations, such as collection stations 120, 125, 130, and 135 illustrated in Figure 1. Specifically, collection stations are "deployed to monitor computer networks within, for example, remote customer sites. In FIG. 1, collection stations 120 and 125 have been deployed to monitor a first computer network 110 (e.g., a customer site designated as "CO"), while collection stations 130 and 135 have been deployed to monitor a second computer network 115 (e.g., a customer site designated as "NY")." (Paragraph [0020].) *Natarajan* explains the problem being addressed as a designated IP address, such as "15.2.112.1", that is used within both computer networks 110 and 115, and as a result, the true source of an event that arises at one device with the duplicate IP address cannot be determined. (Paragraph [0020].)

Natarajan explains that the duplicate IP address issue can be resolved by using at each collection station an "identifier tag that uniquely identifies at least one collection computer monitoring the event," such as having the identifier tag be the name or domain name of the at least one collection computer, having the identifier tag be a customer name as a group name for a collection of computers at a customer site, or having the identifier tag be a unique name for each collection computer residing in the customer site." (Paragraph [0021].)

Note that the focus in *Natarajan* is on uniquely identifying the collection computer, or a group of collection computers, that monitor the events, as opposed to identifying the customer site itself. Also note that because a customer site can have many different domains, the use of the name or domain name of a collection computer or group of collection computers merely identifies a domain within the customer site, not the customer site itself. Even if the identifier tag is taken to be "a customer name," that would be understood to be a name provided or specified by the customer, but such a customer name would not uniquely identify the customer site, only the collection computer or group of collection computers "at a customer site" for which the customer name is used as the identifier tag. Thus, because (1) *Natarajan* defines the "identifier tag" as something that "uniquely identifies at least one collection computer monitoring the event," (2) *Natarajan* does not define the identifier tag as uniquely identifying the customer site, (3) *Natarajan* explains that a domain name for the

collection computer(s), but not a domain name for the customer site, can be used as the identifier tag, and (4) a customer site can have multiple domain names, Natarajan's identifier tag "only uniquely identifies at least one collection computer" as opposed to uniquely identifying the customer site.

Natarajan later explains that the source of an event can be identified to a user by using the identifier tag of the collection computer along with a network address, such as an IP address, of the network element that generated the event, or that the source can be identified using the identifier tag to map the collection computer to a group of collection computers, referred to as a "domain." (Paragraph [0027].) As a result, the management computer can use both (a) the domain name (e.g., "FooNet") that is derived from the identifier tag that uniquely identifies the collection computer and (b) the source name, such as the IP address or domain name of the network node that is down, to uniquely identify the source of the event. (Paragraph [0028].) Thus, neither the domain name nor the source name are sufficient to uniquely identify the source of the event, and therefore neither the domain name nor the source name alone can uniquely identify a customer site. Furthermore, the combination of the domain name and source name only allow *Natarajan* to identify the source of the event, and the source is a particular device. Thus, even once the source of the event is identified, *Natarajan's* approach does not uniquely identify the customer site that includes the particular device, which would presumably be performed by a lookup of the particular device against a collection of information that identifies which devices are included in which customer sites.

Table 400 in *Natarajan* provides an illustrative example. The source 430 for the events 440 and 445 on the sixth and seventh line of the table is given to be the same value, namely "beast.cmd.hp.com" and thus would appear to be identical sources, when in fact the sources are different because they originate in different domains. (Paragraph [0029].) Specifically, the management station can determine the true source of both events 440 and 445 on the sixth and seventh lines of Table 400, respectively, based on the domains that are different for each event as indicated in the station field 425 of Table 400, namely "hpcndsn.cnd.hp.com" and "nityant.cnd.hp.com", respectively. (Paragraph [0029].)

Also, *Natarajan* explains that either the management computer, such as management station 104, or a collection computer can perform hostname resolution of objects managed by each collection computer. For the latter, the host name resolved by a collection computer can

be included in the information sent to the management computer, thereby preventing the management computer from failing due to an inability to resolve a hostname or due to duplicate hostnames. (Paragraph [0030].)

Finally, note that as indicated in Figure 1 the collection stations for each network are separate from the network elements that the collection stations are monitoring. For example, collection stations 120 and 125 that monitor network 110 are separate from network element 140, and similarly, collection stations 130 and 135 that monitor network 115 are separate from network element 145. Note that Figure 1 is illustrating the problem that arises from both of network elements using the same unregistered IP address of 15.2.112.1, which is unique within each of networks 110 and 115, but which is obviously not unique among networks 110 and 115 taken together since each of networks 110 and 115 is using that same IP address for different network elements. Thus, when an event is generated from either of network elements 140 and 145 and one of the collection stations for each network 110 and 115 forwards the event to management station 105, management station 105 cannot determine from which of networks 110 or 115 the event occurred because the IP address seen by management station is used by two different network elements. However, by including the identifier tag that uniquely identifies the collection station or collection stations for each network, along with the duplicate IP address, the management station 105 is able to identify the source of an event as being either network element 140 in network 110 or network element 145 in network 115.

(4) THE OFFICE ACTION'S CITATIONS TO *NATARAJAN* IN REJECTING CLAIM 1

The Office Action states that "Natarajan shows wherein the device is associated with a particular site in a plurality of sites and said network operations center processes alarms for each site in the plurality of sites (Figs 1, 2, 4A, [0029-0030])."

However, Claim 1 as amended above, features "the network device detecting an event within the network device," "the network device propagating an alarm to an alarm identification component that is hosted within the network device," and "the alarm identification component augmenting the alarm with identification information resulting in creating an augmented alarm." Thus, the network device both (a) detects the event and

(b) hosts the alarm identification component that augments the alarm with the identification information.

In contrast to the approach of Claim 1, *Natarajan* discloses that an event at a network element, such as network element 140 of network 110 or network element 145 of network 115, as illustrated in Figure 1, is detected through the monitoring of the networks by the collection stations, namely collection stations 120 and 125 for network 110 and collection stations 130 and 135 for network 115. *Natarajan* explains that collection stations send either the name or the IP address of the network element that generated the event as the source of the event, but that when two events come from different collection stations that monitor different network elements that have the same IP address, the management station is unable to correctly identify the source of a particular event as having come from one of the two possible sources. (Paragraph [0009].)

Natarajan's approach to solve this problem is to associate an identifier tag with an event in which "the identifier tag ***uniquely identifies*** at least one **collection computer** monitoring the event," so that when the management computer receives the information about the event from the collection computer, the "management computer ***derives*** **an identification of each collection computer** from the identifier tag" so that the "source of the event is identified to a user using the identification of each collection computer."

Thus, in *Natarajan's* approach, the identifier tag is not originally included in the information about the event, but rather is added by the collection computer, as illustrated in Figure 3 of *Natarajan*. Specifically, Figure 3 illustrates PDU 310 that is an example of how an event was previously reported by a collection station (paragraph [0023]), along with PDU 312 that is an example of how an event is reported by including an additional variable binding 304, such as "foonet.com" that is the name or domain name for the collection computer and thus is used as the identifier tag that uniquely identifies that collection computer or group of collection computers monitoring the event (paragraph [0024]).

In contrast, the "identification information that uniquely identifies the particular site among the plurality of sites" in the approach of Claim 1 is used to augment the alarm by the alarm identification component that is hosted by the network device within which the event occurs. Thus, in the approach of Claim 1, the augmenting of the alarm with the identification information is not done by some other network device, but rather is performed at the same

network device that detects the event that occurs within that same network device. This is different than in *Natarajan* in which an event occurs at a network element, such as network element 140 or 145, and then a collection station, such as one of collection stations 120, 125, 130, and 130, includes the additional variable binding that uniquely identifies the particular collection station that is sending the information to the management station about the event at the network element that the collection station is monitoring.

In particular, one of the examples of an "event" that *Natarajan* provides is that a node is down (see Figures 4A and 4B). A node that is "down" would be unable to detect that event, even though the event occurs at the node, because of the node being down. Yet despite the node being down, the event is still detected and reported to the management station by the collection station monitoring the node. Thus, *Natarajan's* approach does not involve the network element at which the event occurs sending an alarm, little less an augmented alarm or an augmented alarm based on "identification information that uniquely identifies the particular site." Rather, in *Natarajan's* approach, the information sent about the event at the node is from the collection station, not the node at which the event occurs.

Natarajan is consistent in referring to the identifier tag as "uniquely identifying at least one collection computer" that is monitoring the event. (See Abstract, paragraphs [0012], [0018-0019], [0021], [0024], and Claims 1 and 7 of *Natarajan*.) Thus, the identifier tag as described by *Natarajan* does not uniquely identify a customer site from among other customer sites.

Furthermore, upon receipt of the information about the event from the collection computer, the management station must ***derive*** the identity of the collection computer from the identifier tag. (See Abstract, paragraphs [0012] and [0026], and Claims 1, 3, and 7 of *Natarajan*). Therefore, because the management station is deriving the identity of the collection computer from the identifier tag, the identifier tag does not uniquely identify the customer site within which the collection computer resides.

Thus, the Applicant respectfully submits that *Tse* fails to disclose, teach, suggest, or render obvious "the network device detecting an event within the network device," "the network device propagating an alarm to an alarm identification component that is hosted within the network device," "the alarm identification component augmenting the alarm with identification information resulting in creating an augmented alarm," and "identification

information uniquely identifies the particular site among the plurality of sites," as featured in Claim 1.

(5) CONCLUSION OF DISCUSSION OF CLAIM 1, *TSE*, AND *NATARAJAN*

Because *Tse* and *Natarajan*, either alone or in combination, fail to disclose, teach, suggest, or render obvious "the network device detecting an event within the network device," "the network device propagating an alarm to an alarm identification component that is hosted within the network device," "the alarm identification component augmenting the alarm with identification information resulting in creating an augmented alarm," and "identification information uniquely identifies the particular site among the plurality of sites," as featured in Claim 1, the Applicant respectfully submits that Claim 1 is allowable over the art of record and is in condition for allowance

C. CLAIMS 22, 43, AND 64

Claims 22, 43, and 64 contain features that are either the same as or similar to those described above with respect to Claim 1, although in the context of a computer-readable medium, a system with means elements, and a system without means elements, respectively. In particular, each of Claims 22, 43, and 64 feature "the network device detecting an event within the network device," "the network device propagating an alarm to an alarm identification component/means that is hosted within the network device," "the alarm identification component/means augmenting the alarm with identification information resulting in creating an augmented alarm," and "identification information uniquely identifies the particular site among the plurality of sites," which are the same as in Claim 1. Therefore, based on at least the reasons stated above with respect to Claim 1, the Applicant respectfully submits that Claims 22, 43, and 64 are allowable over the art of record and are in condition for allowance.

D. CLAIMS 2-21, 23-42, 44-64, AND 65-84

Claims 2-21 depend upon Claim 1, Claims 23-42 depend upon Claim 22, Claims 44-64 depend upon Claim 43, and Claims 65-84 depend upon Claim 64, and thus include each and every feature of the corresponding independent claims. Each of

Claims 2-21, 23-42, 44-64, and 65-84 is therefore allowable for the reasons given above for Claims 1, 22, 43, and 64. In addition, each of Claims 2-21, 23-42, 44-64, and 65-84 introduces one or more additional limitations that independently render it patentable. However, due to the fundamental differences already identified, to expedite the positive resolution of this case a separate discussion of those limitations is not included at this time. Therefore, it is respectfully submitted that Claims 2-21, 23-42, 44-64, and 65-84are allowable for the reasons given above with respect to Claims 1, 22, 43, and 64.

CONCLUSION

The Applicant believes that all issues raised in the Office Action have been addressed and that allowance of the pending claims is appropriate. After entry of the amendments, further examination on the merits is respectfully requested.

The Examiner is respectfully requested to contact the undersigned by telephone if it is believed that such contact would further the examination of the present application.

For the reasons set forth above, it is respectfully submitted that all of the pending claims are now in condition for allowance. Therefore, the issuance of a formal Notice of Allowance is believed next in order, and that action is most earnestly solicited.

To the extent necessary to make this reply timely filed, the Applicant petitions for an extension of time under 37 C.F.R. § 1.136.

If any applicable fee is missing or insufficient, throughout the pendency of this application, the Commissioner is hereby authorized to any applicable fees and to credit any overpayments to our Deposit Account No. 50-1302.

Respectfully submitted,

HICKMAN PALERMO TRUONG & BECKER LLP

Date: October 24, 2007

/CraigGHolmes#44770/

Craig G. Holmes
Reg. No. 44,770

2055 Gateway Place, Suite 550
San Jose, CA 95110-1089
Telephone: (408) 414-1207
Facsimile: (408) 414-1076